Optimal Mixtures of Different Types of Recovery Schemes in Optical Networks

David Griffith, Kotikalapudi Sriram, Stephan Klink, and Nada Golmie
National Institute of Standards and Technology (NIST)
100 Bureau Drive, Stop 8920
Gaithersburg, MD 20899-8920
Email: david.griffith@nist.gov

Abstract

A network operator who wants to provide recovery services to end-users has a range of options from which to choose. There is a fundamental trade off between recovery speed and robustness on one hand and the amount of resource overhead that is consumed by backup paths on the other hand. Several previous studies have examined this trade-off qualitatively and quantitatively. This paper provides a valuable extension to the previous work by considering the behavior of networks that support multiple types of protection schemes. Using simulations, we examined a network that supported both 1+1 and 0:1 recovery and determined the optimal mixture of the two schemes as a function of network load and of various performance metrics such as restorability of failed connections.

I. Introduction

Optical communications networks have evolved in recent years from systems that carry only voice traffic for telephony to data transport networks that support a variety of applications and traffic types [1]. Because of the convergence of voice, video, and data services that has taken place in the last decade, service providers' networks now carry vast quantities of critical, delay-sensitive, loss-sensitive traffic. In such networks, it is important for service providers to be able to offer a range of services and capabilities to their customers. One of the most important services is rapid failure recovery. Failures that result from a disruption of the physical network can have serious economic consequences, particularly if the affected traffic cannot be recovered for a long period of time. Legacy carrier networks employed a variety of failure recovery mechanisms, such as redundant connections in the control plane of networks using Signaling System 7 (SS7) and Automatic Protection Switching (APS) in Synchronous Optical Network (SONET) networks. With the transition to optical switching networks with mesh connectivity, it has become necessary to introduce path protection techniques, in which traffic on a failed connection is automatically shunted onto a protection, or backup, path within a short period of time.

There have been many discussions of path protection architectures for next generation optical networks in the literature, such as [2], [3], [4], [5], [6], and [7]. In addition, the Internet Engineering Task Force (IETF), through the Common Control and Measurement Plane (CCAMP) working group, generated a set of documents [8], [9], and [10] that are intended to describe the protection and restoration capabilities that are supported by the Generalized Multi-Protocol Label Switched (GMPLS) control plane for optical networks and to define the signaling to carry out both path and span protection. Part of this description is a classification system for recovery schemes that uses the degree of backup resource pre-provisioning as the primary differentiator [9]. This system is accompanied by a qualitative discussion that describes some of the design trade-offs associated with each of the recovery schemes. Last year, the authors supported the development of this document by carrying out a quantitative, simulation-based analysis of the recovery schemes described in the analysis draft [11]. That analysis assumed that only one

type of recovery scheme was being used in the network. It is more likely, however, that a network operator will want to offer different levels of protection to prospective customers, e.g., dedicated backup paths as a form of premium service. As a result, we can expect that there will be a mixture of different types of recovery schemes, with various associated levels of pre-provisioning, in use at the same time. The goal of this paper is to simulate the performance of a network that uses more than one recovery scheme and to determine what mixture of usage of the different schemes yields the best network performance.

In this paper, we restrict the scope of our analysis to the following two recovery schemes: 1+1 protection and dynamic recovery (sometimes also known as 0:1 protection). The former scheme protects traffic by creating two lightpaths and transmitting duplicate copies of the data on both paths. The latter scheme provides no protection resources for the data traffic prior to a failure event. Instead, if a lightpath with 0:1 protection is disrupted, the traffic that it is carrying is shifted to a new lightpath that is created on the fly. In some cases, it may not be possible to create a new lightpath, in which case the traffic is lost. The two schemes that we consider represent two extremes with respect to the degree of resource pre-provisioning that they require. There exist other approaches that lie somewhere between these two, described in [9] and [11], and this analysis can be extended to include these.

The remainder of this paper is organized as follows. In Section II, we describe the simulations that we performed to determine the optimal mixture of the two recovery schemes over a range of traffic loads. In Section III, we present our experimental results and also show that the optimal mixture of heterogeneous recovery schemes depends on the relative importance of the various design criteria that the network operator must satisfy. In Section IV, we summarize our results and discuss how our work may be extended to examine other types of protection schemes.

II. Simulation Description

In this section we describe the simulations that we used to determine the optimal mixture of recovery schemes in a wide area network (WAN). We extended the set of simulations that we performed in [11] using the GMPLS/Lightwave Agile Switching Simulator (GLASS) simulation tool [12], which is based on the SSF framework [13]. Using the NSFNet topology, shown in Fig. 1, we loaded the network with connections whose arrival times follow a Poisson process with rate parameter λ and whose durations are exponentially distributed with mean $1/\mu$.

The source and destination node IDs were uniformly distributed over the entire node ID space, with the restriction that calls could not originate and terminate at the same node. The probability that a given connection request was protected using dynamic recovery, i.e., 0:1 protection, was set to p, so that the arrival rate for 0:1 connections was $p\lambda$. Thus the probability that a new connection request received 1+1 protection was $(1-p)$, and the associated arrival rate was $(1-p)\lambda$. The connection duration was exponentially distributed with a mean value of $1/\mu = 100$ hours. The average arrival rate of the connection requests and the duration of the connections can be used to calculate the load offered to the network. For example, given μ, if $\lambda = 0.01$ requests/hour (corresponding to an average gap of approximately four days between requests), then the network sees a load of one Erlang. In our simulations, we considered loads from 100 to 1000 Erlangs, at intervals of 100 Erlangs;

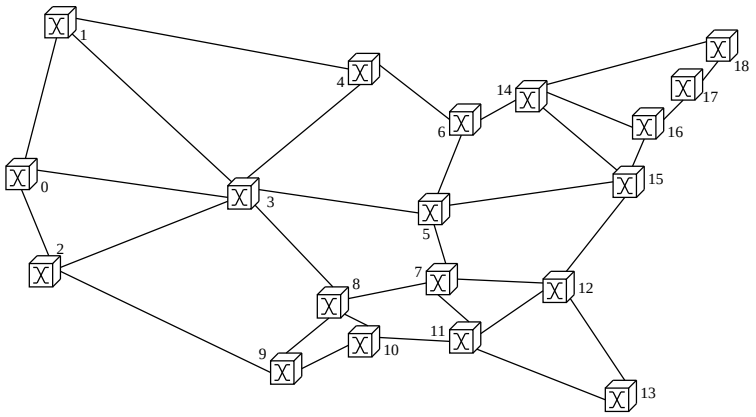

Fig. 1. The NSF network topology.

the upper and lower limits correspond to mean interarrival times between connection requests of two hours and six minutes, respectively.

Each connection request was randomly assigned to have either 1+1 protection (with probability $1 - p$) or 0:1 protection with probability p. A new request for a 1+1-protected connection requires two edge-disjoint paths between the designated source and destination nodes, whereas a 0:1-protected connection requires only a working path to be set up. In both cases, we needed to find a suitable working path to carry the requested traffic flow. To do this, we implemented a modified version of the K-Shortest-Path algorithm described by Bhandari in [14]. We used the algorithm to compute a maximum of 5 edge-disjoint paths between the connection endpoints. In some cases, it was not possible to obtain 5 edge-disjoint paths; the network simply used the smaller number of paths that were available when this occurred. The links were weighted by distance, since it is desirable to have the shortest possible optical path across the network. Starting with the best path, the network controller attempted to find a wavelength that could be assigned to the path. It used the FirstFit algorithm to do this; if a wavelength could be found, then the path was designated as the working path. If the wavelength assignment algorithm was unsuccessful, the network controller would move on to the next best path. If the supply of candidate paths was exhausted, the connection request was considered failed. If the requested connection used dynamic protection, then the connection setup was completed at this point. If the requested connection was 1+1-protected, and if a working path could be found, the network controller would begin to set up the backup path. The controller did this by examining the remaining members of the set of candidate paths and using a modified wavelength assignment algorithm, called FirstFitBackup. This algorithm is similar to FirstFit, but it allows multiple backup paths that overlap on one or more links to share a single wavelength on those links. Different backup paths were allowed to share resources only if their respective working paths were link-disjoint. If only the working path was available for a 1+1-protected connection, then the connection was blocked.

To generate the performance statistics reported in Section III, we allowed the network to reach steady-state and then failed a single node. We implemented a node failure by simultaneously causing all the unidirectional fibers originating from the affected node to fail. We simulated each node's failure in three different runs. We next formed a network-wide ensemble average by

repeating the experiment for each node in the network and averaging the final value of each metric over the full set of runs. Using the data from the 54 runs that we carried out, we computed three metrics, P_B, R, and U, which we describe below.

The new connection blocking probability, P_B, is obtained from the number of new connection requests that cannot be satisfied because the required network resources are not available. A connection request fails if it is not possible to successfully complete both the working path setup and the backup path pre-provisioning setup steps described above. To compute P_B, we maintained two counters that were updated upon the arrival of each new connection request. One counter kept track of the total number of new connection requests, and was incremented upon the arrival of each new request, while the other tracked the number of failed requests and was incremented only when a request failed. The estimated value of P_B is the ratio of the second counter to the first. The restorability, R, is the probability, expressed as a percentage, that a failed connection can be successfully recovered, and is a measure of the network's resilience to failures. When we computed this metric for each node failure scenario, we excluded connections that terminate at the failed node, as their being restored was impossible. We computed this metric by tallying all the connections that were broken by a node failure event, as well as all those that successfully switched to their backup paths. The measured value of R was the ratio of the number of recovered connections to the total number of broken connections. The link utilization, U, is a measure of the amount of traffic that the network supports. It is expressed as the percentage of link bandwidth that is occupied by active connections, averaged over all links on the network and also averaged over the duration of the simulation run. Both working and backup paths for 1+1 connections contribute to the total utilization.

III. Experimental Results

In this section, we describe the results that we obtained from the set of simulations that we discussed in Section II. We plot each of the three metrics, P_B, R, and U, with respect to the network load and the fraction of new connections that use dynamic recovery rather than 1+1 protection. We also show how attempting to simultaneously optimize a subset of the three metrics produces a set of optimal values for p that can be expressed as a function of the network load and the relative importance of two components of the composite performance metric that we are trying to minimize.

We plot the new connection blocking probability, P_B, as a function of the network load in Fig. 2 for various values of the fraction of 0:1 connections, p. In Fig. 3, we plot the blocking probability versus p for various values of load. The graphs show that the new connection blocking probability decreases monotonically as p increases, for any given load. Furthermore, P_B increases with load for any value of p. This follows from the fact that a lower value of p corresponds to a larger quantity of dedicated backup paths that consume the available bandwidth on each of the links in the network. Thus the price for a greater degree of robustness is a significant increase in overhead that directly affects the ability of the network to support additional connections, except at very low loads (200 Erlangs or less).

In Fig. 4 and Fig. 5 we plot the restorability, R, with respect to the network load and with respect to p. We observe the same general trends that were associated with P_B in Fig. 2 and Fig. 3. For any given load, increasing the value of p produces a decrease in R, in the sense that $R(1) < R(0)$. However, were able to observe a monotonic decrease in R only for certain loads,

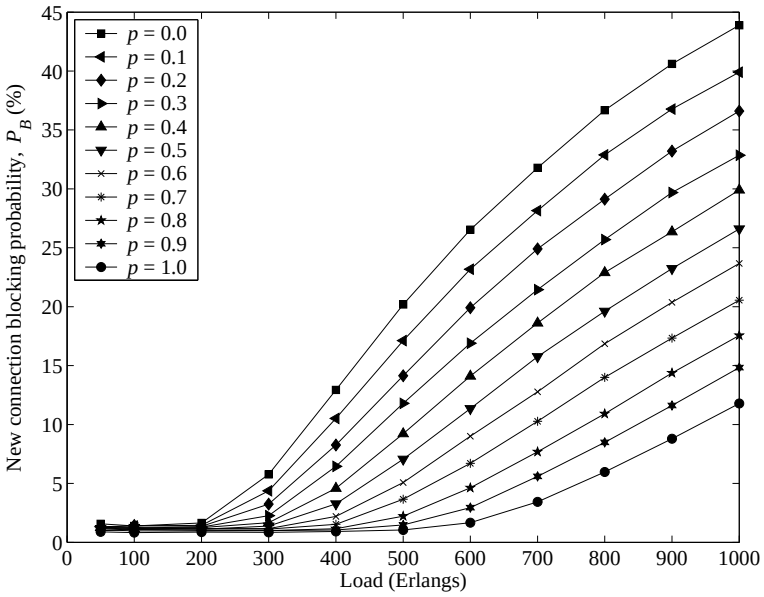

Fig. 2. Plot of new connection blocking probability, P_B, versus network load.

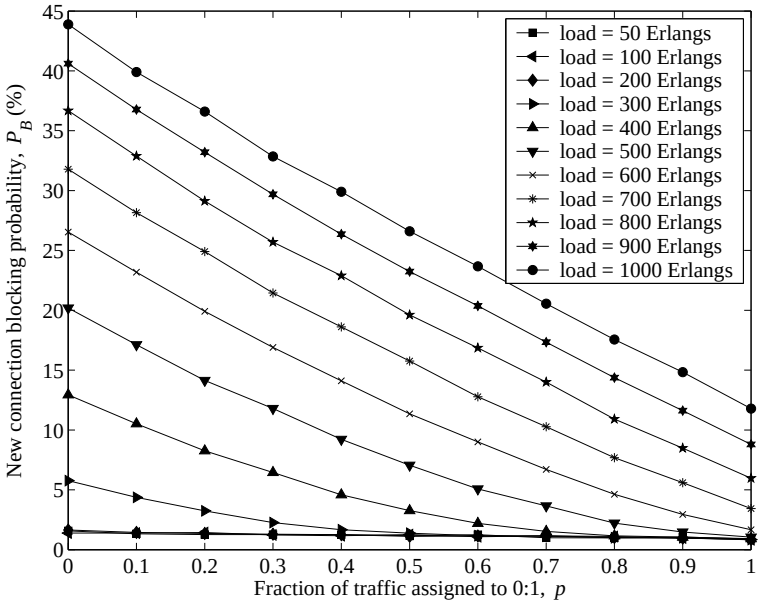

Fig. 3. Plot of new connection blocking probability, P_B, versus faction of 0:1 traffic, p.

e.g., 700 Erlangs. We believe that much of the variation in R, especially at lower loads, is due to noise effects. The decrease in R with respect to p is more pronounced at higher loads. This follows from the fact that a higher value of p corresponds to a greater proportion of 0:1 connections in the network. The overall restorability can become quite low, especially at high loads, if we are forced to rely exclusively on dynamic recovery when free bandwidth on an alternate path may not be available. If we increase the proportion of 1+1-protected connections, we naturally gain an improvement in restorability since a 1+1 protected connection can be lost only if both component paths are disabled by a single failure event. This can occur when both paths pass through a common node, since the two paths that support a 1+1-protected traffic flow are required to be link-disjoint but

not necessarily node-disjoint.

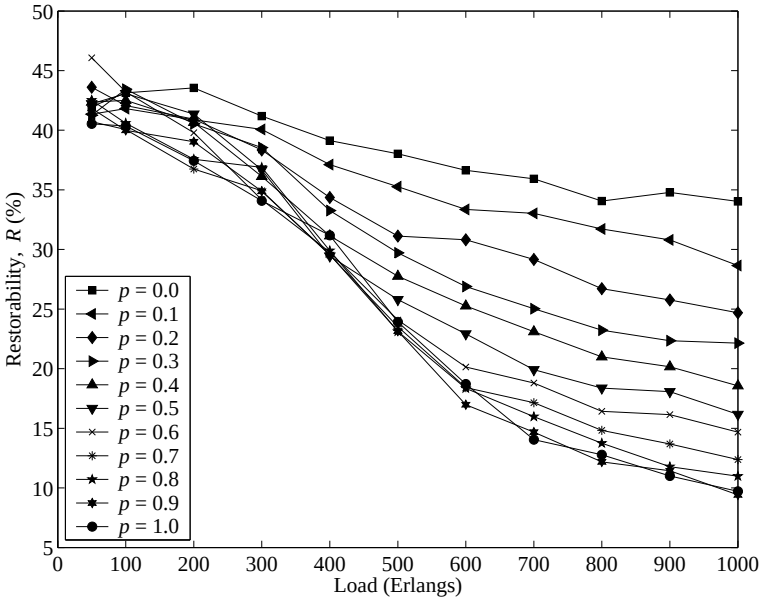

Fig. 4. Plot of restorability, R, versus network load.

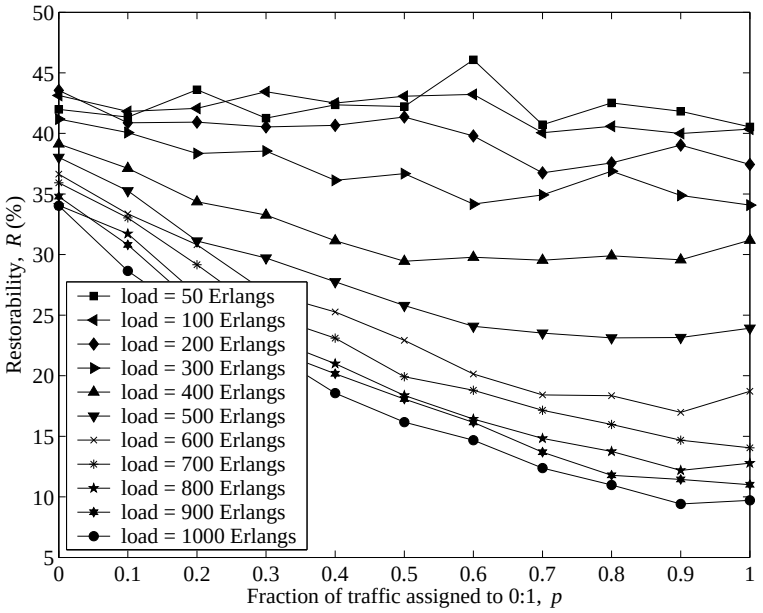

Fig. 5. Plot of restorability, R, versus faction of 0:1 traffic, p.

Finally, we plot the link utilization, U, versus load and p in Fig. 6 and Fig. 7, respectively. We observe the same behavior with respect to both parameters as occurs in Fig. 2 and Fig. 3. The average link usage is greatest, for a given load, when we use 1+1 protection. This results naturally from the requirement to have two diverse paths for every connection. In addition, for a given value of p, we see the highest link usage when the load is largest. Interestingly, there is relatively little variation in U as p varies from 0 to 1, and at very low loads the U-curve is nearly flat, as Fig. 7 shows. We observe the greatest variation with respect to p for loads in the range of 400 to 500 Erlangs.

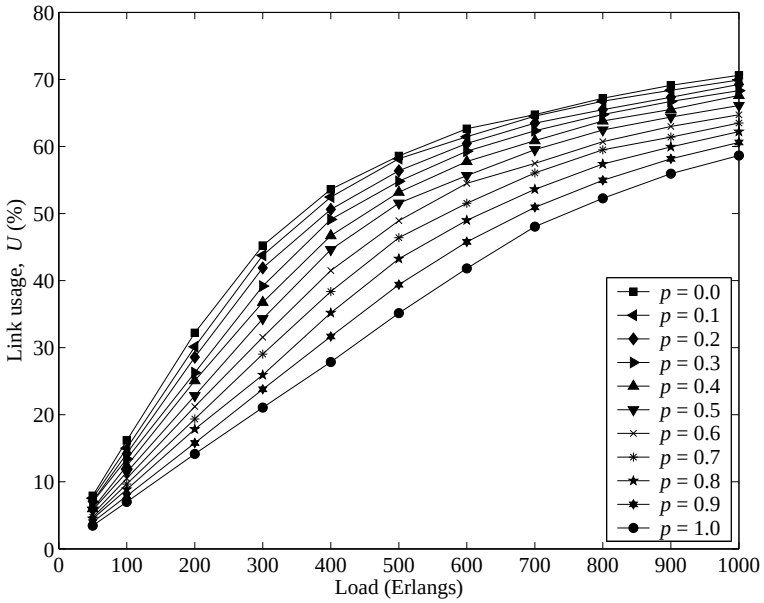

Fig. 6. Plot of link utilization, U, versus network load.

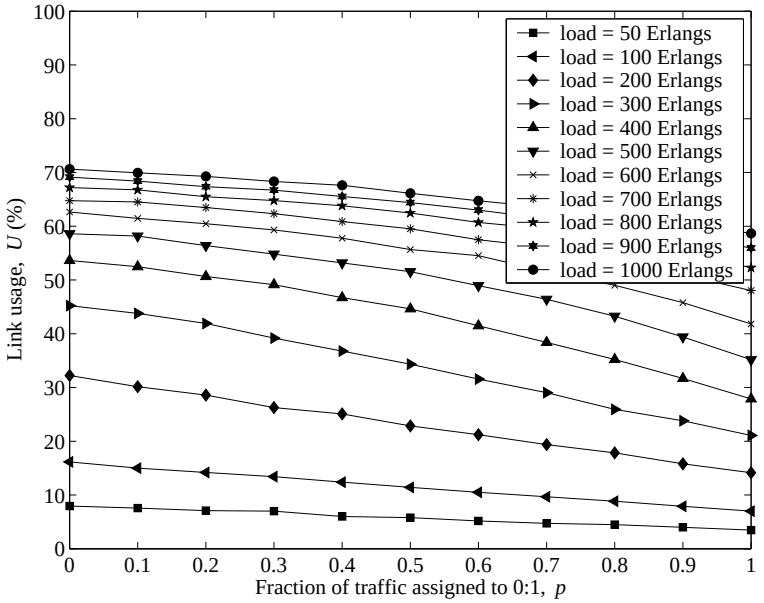

Fig. 7. Plot of link utilization, U, versus faction of 0:1 traffic, p.

If we consider any of the above metrics in isolation, we find that the optimal solution for any given load is to use only one type of protection scheme for all the connections that the network must support. For instance, if we are interested in minimizing link utilization, we should use dynamic restoration for every connection. Conversely, if maximizing restorability is our primary goal, then every new connection should have a dedicated backup path assigned to it. More often, our goal is to optimize the performance of the network with respect to a combination of performance metrics, which sometimes cannot all

be minimized simultaneously. To examine this issue, we define a composite performance metric of the form

$$J(p) = \alpha Q_1(p) + (1-\alpha)Q_2(p), \tag{1}$$

where $Q_1(p)$ and $Q_2(p)$ are two cost functions that we want to minimize and α is an importance parameter that allows us to tune the composite metric by emphasizing one cost function over the other. We can find a global minimum for $J(p)$ at $p = p^*$ if

$$\alpha \frac{dQ_1(p)}{dp}\bigg|_{p=p^*} = (\alpha - 1)\frac{dQ_2(p)}{dp}\bigg|_{p=p^*}. \tag{2}$$

In general, we can use the data associated with the three metrics of interest (new connection blocking probability, restorability, and link utilization) to create three composite metrics that we can then optimize with respect to p. We begin by defining the metric

$$J(p) = \alpha P_B(p) + (1-\alpha)(100 - R(p)), \tag{3}$$

which would be used by an operator who is interested in achieving a balance between maximizing the connection restorability and minimizing the new connection blocking probability. For example, setting α equal to one gives a network performance metric that is minimized if we minimize the new connection blocking probability, $Q_1(p) = P_B(p)$. We set $Q_2(p) = 100 - R(p)$, which is the percentage of connections that could not be restored, because we are interested in maximizing the restorability rather than minimizing it.

Using the data shown in Fig. 2, Fig. 3, Fig. 4, and Fig. 5, we constructed the performance metric $J(p)$ over a range of values $0 \leq \alpha \leq 1$. For each value of α that we considered, we obtained the value of p that minimized $J(p)$ for a given load. The resulting set of optimal p values is plotted in Fig. 8 as a solid mesh superimposed over a contour plot that lies in the plane $p^* = -1.5$. In this figure and the following two figures, the three contour lines are associated with the following set of values of p^*, which are superimposed on the corresponding lines: $\{0.1, 0.5, 1\}$. We immediately notice the abrupt discontinuity at $\alpha \approx 0.4$ for loads greater than 400 Erlangs, indicating that a homogeneous protection arrangement for the entire network is appropriate at high loads. The challenge facing the network operator is to determine the relative importance of the robustness and new connection blocking requirements, which is critical given the abrupt nature of the discontinuity. At lower loads, we should use dynamic recovery exclusively (i.e., set $p = 1$) only if α is close to 1. For other values of α, a mixture of recovery types is better. For instance, when $\alpha = 0.6$, $p^* \approx 0.5$ at very low loads. As the load increases, p^* initially decreases to less than 0.1 before increasing gradually to 1.0. We also note that there is a more gradual transition as the load varies, for a fixed value of α. Further investigation is required to determine a cause for the different types of transitions seen in the figure.

Next we consider a composite metric formed from the new connection blocking probability and the link utilization. In this case, our objective is to minimize the new connection blocking probability while maximizing the average link utilization. By maximizing link utilization, we are insuring that a large quantity of protection bandwidth is available. The resulting cost

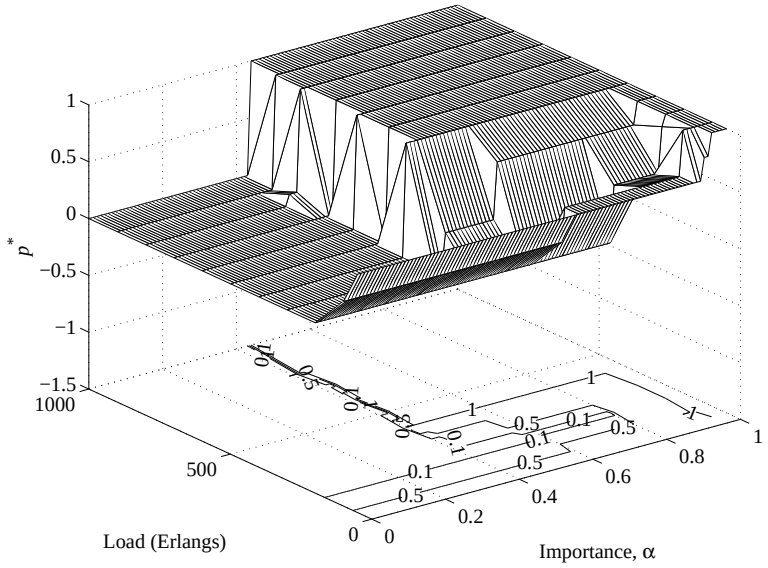

Fig. 8. Mesh and contour plots of p^* versus network load and α, where we are trying to minimize the new connection blocking probability while maximizing restorability.

function is

$$J(p) = \alpha P_B(p) + (1 - \alpha)(100 - U(p)). \tag{4}$$

The resulting surface, shown in Fig. 9, does not have the sharp boundary that we observed in Fig. 8. Instead, the (load, α) plane is divided into three regions that are clearly visible in the figure. If the load is high and $\alpha > 0.4$, only dynamic recovery should be used. Likewise, if the load is low or if α is small (less than 0.1), then only 1+1 connections should be set up. There is a large transitional region in the middle of the plane where p^* is relatively sensitive to changes in the two parameters. For instance, Fig. 9 indicates that approximately half of new connection arrivals should be assigned a dedicated backup path if the network load is 500 Erlangs and $0.4 < \alpha < 0.55$, i.e., if the new connection blocking probability and the link utilization are of roughly equal importance.

Finally, we consider the case where the composite performance metric is based on the link utilization and the restorability. The metric is

$$J(p) = \alpha U(p) + (1 - \alpha)(100 - R(p)), \tag{5}$$

since we want to maximize the probability that a failed connection is restored while also minimizing the number of connections on each link. As before, we generated the surface shown in Fig. 10 by computing $J(p)$ using multiple values of α and then locating of the value of p associated with the minimum value of $J(p)$ for each value of the network load. In this case, we have another abrupt transition, similar to the one in Fig. 8, stretching diagonally across the (load, α) plane. We expect to see a degree of similarity between Fig. 8 and Fig. 10 because of the strong correlation between the new connection blocking probability and the link utilization. The relative importance of the two cost functions in equation (5) as the trigger for determining which

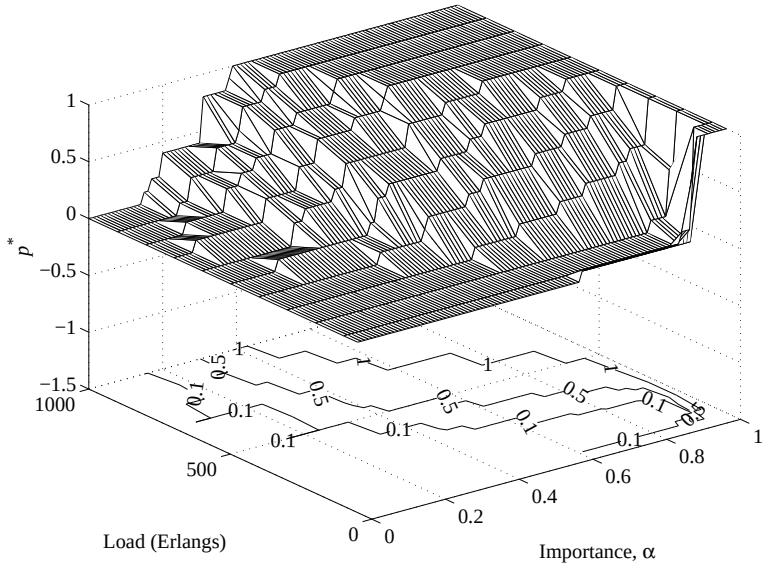

Fig. 9. Mesh and contour plots of p^* versus network load and α, where we are trying to minimize the new connection blocking probability while maximizing link utilization.

protection scheme to use can be seen to vary with respect to the network load. This has very interesting implications for the network operator, since using a particular protection scheme for every connection to the network can become a suboptimal solution once the network load increases beyond a certain point.

Because of some of the issues raised in the preceding discussion, such as the abrupt transitions that make the choice of recovery strategy very sensitive to network conditions and management constraints, and may be more appropriate to employ a range of recovery strategies, such as those using four different levels of pre-provisioning as described in [9]. We are planning to conduct additional studies using these different recovery schemes to determine the best proportions of each that should be used to optimize the network performance with respect to the metrics that we discussed above. The simulations will also determine whether lower sensitivity to variations in load and α can be achieved by using a larger set of the recovery techniques.

IV. Summary

In this paper, we described the result of the set of simulations that we performed to determine the optimal mixture of a heterogeneous set of failure recovery schemes. We demonstrated that the optimal ratio of the different recovery schemes depends strongly on which criteria the network operator is using. If restorability, i.e., the probability that a failed connection can be successfully recovered, is the primary design concern, then the majority of connection requests should be set up with 1+1 protection. If, however, minimizing blocking of new connections is more important, then the greater proportion of the lightpaths should be set up using dynamic recovery, i.e., 0:1 protection. Often, the optimal network design involves achieving a balance between these competing goals. The results of our simulations indicate that if the relative importance of these design goals is known, perhaps based on the network operators' insight and experience, then an optimal mixture of recovery schemes can be chosen for a given network load.

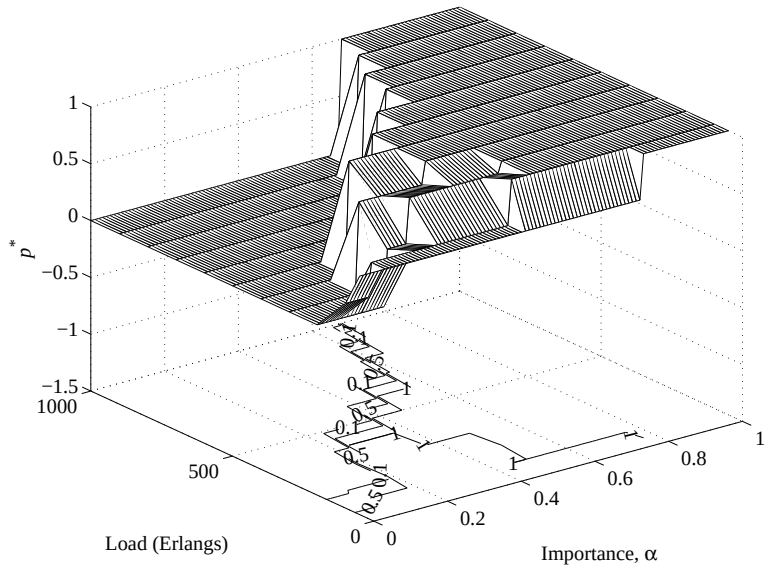

Fig. 10. Mesh and contour plots of p^* versus network load and α, where we are trying to minimize link utilization while maximizing restorability.

REFERENCES

[1] ITU-T Recommendation G.8080/Y.1304, "Architecture of the automatically switched optical network (ASON)" (International Telecommunication Union-Telecommunication Standardization Sector, November 2001), http://www.itu.int/ITU-T/

[2] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM mesh networks," *Journal of Lightwave Technology*, vol. 21, no. 4, pp. 870–883, April 2003.

[3] R. R. Iraschko and W. D. Grover, "A highly efficient path-restoration protocol for management of optical network transport integrity," *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 779-794, Oct. 2000.

[4] G. Li, J. Yates, R. Doverspike, and D. Wang, "Experiments in fast restoration using GMPLS in optical/electronic mesh networks," *Proceedings of the Optical Fiber Communication Conference and Exhibit, 2001*, vol. 4, pp. PD34-(1–3).

[5] M. Goyal, G. Li, and J. Yates, "Shared mesh restoration: a simulation study," *Proceedings of the Optical Fiber Communication Conference and Exhibit, 2002*, pp. 489–490.

[6] Y. Xiong, D. Xu, and C. Qiao, "Achieving fast and bandwidth-efficient shared-path protection," *Journal of Lightwave Technology*, vol. 21, no. 2, pp. 365–371, Feb. 2003.

[7] Qin Zheng and G. Mohan, "Protection approaches for dynamic traffic in IP/MPLS-over-WDM networks," *IEEE Communications Magazine*, vol. 41, no. 5, pp. S24–S29, May 2003.

[8] "Recovery (Protection and Restoration) Terminology for GMPLS," E. Mannie and D. Papadimitriou, Eds., IETF Internet draft.

[9] "Analysis of Generalized MPLS-based Recovery Mechanisms (including Protection and Restoration)," D. Papadimitriou and E. Mannie, Eds., IETF Internet draft.

[10] "Generalized MPLS Recovery Functional Specification," J. Lang and B. Rajagopalan, Eds., IETF Internet draft.

[11] D. Griffith, R. Rouil, S. Klink, and K. Sriram, "An Analysis of Path Recovery Schemes in GMPLS Optical Networks with Various Levels of Pre-Provisioning," in Proceedings of SPIE Vol. 5285 *OptiComm 2003: Optical Networking and Communications*, edited by Arun K. Somani, Zhensheng Zhang, (SPIE, Bellingham, WA, 2003), pp. 197–208.

[12] Oliver Borchert, Richard Rouil, "The GMPLS Lightwave Agile Switching Simulator - An overview," www.antd.nist.gov/glass

[13] *Scalable Simulation Framework API Reference Manual, Version 1.0*, James H. Cowie, Ed., www.ssfnet.org/SSFdocs/ssfapiManual.pdf.

[14] R. Bhandari, "Optimal physical diversity algorithms and survivable networks," *Proceedings of the Second IEEE Symposium on Computers and Communications, 1997*, Alexandria, Egypt, 1-3 July, 1997, pp. 433-441.